

TRANSACTIONS

Example: Transfer Euro 50 from A to B



The image shows a SEPA Euro-Überweisung form. The form is titled 'Euro-Überweisung' and 'Kreditinstitut Überall'. It contains fields for 'Angehöriger zum Begünstigten', 'IBAN', 'Kontonummer des Empfängers', 'Betrag Euro, Cent', 'Kontos des Überweisenden', and 'Angehöriger zum Kreditinstitut'. The form is marked with 'SEPA' on the right side.

1. Read balance of A from DB into Variable a : **read**(A, a);
2. Subtract 50.- Euro from the balance: $a := a - 50$;
3. Write new balance back into DB: **write**(A, a);
4. Read balance of B from DB into Variable b : **read**(B, b);
5. Add 50,- Euro to balance: $b := b + 50$;
6. Write new balance back into DB: **write**(B, b);

TRANSACTIONS

Definition: Transaction

Sequence of DML/DDDL statements

Transforms the data base from one consistent state to another consistent state

ACID-Principle

Transactions obey the following four properties

- **Atomicity:** "All or Nothing"-Property (error isolation)
 - Undo changes if there is a problem
- **Consistency:** Maintaining DB consistency (defined integrity constraints)
 - Check integrity constraints at the end of a TA
- **Isolation:** Execution as if it is the only transaction in the system (no impact on other parallel transactions)
 - Synchronize operations of concurrent TAs
- **Durability:** Holding all committed updates even if the system fails or restarts (persistency)
 - Redo changes if there is a problem

Types of Failures: R1-R4 Recovery

1. Abort of a single TA (application, system)
 - *R1* Recovery: Undo a single TA
2. System crash: lose main memory, keep disk
 - *R2* Recovery: Redo committed TAs
 - *R3* Recovery: Undo active TAs
2. System crash with loss of disks
 - *R4* Recovery: Read backup of DB from tape

ACID-Principle cont.

System guarantees the ACID properties

→ Task of the application programmer?

Define borders of transactions

- as large as necessary
- as small as possible

Programming with Transactions

- **begin of transaction (BOT)**: Starts a new TA
- **commit**: End a TA (success).
 - Application wants to make all changes durable.
- **abort**: End a TA (failure).
 - Application wants to undo all changes.
- N.B. Many APIs (e.g., JDBC) have an auto-commit option:
 - Every SQL statement run in its own TA.

SQL Example

insert into Lectures

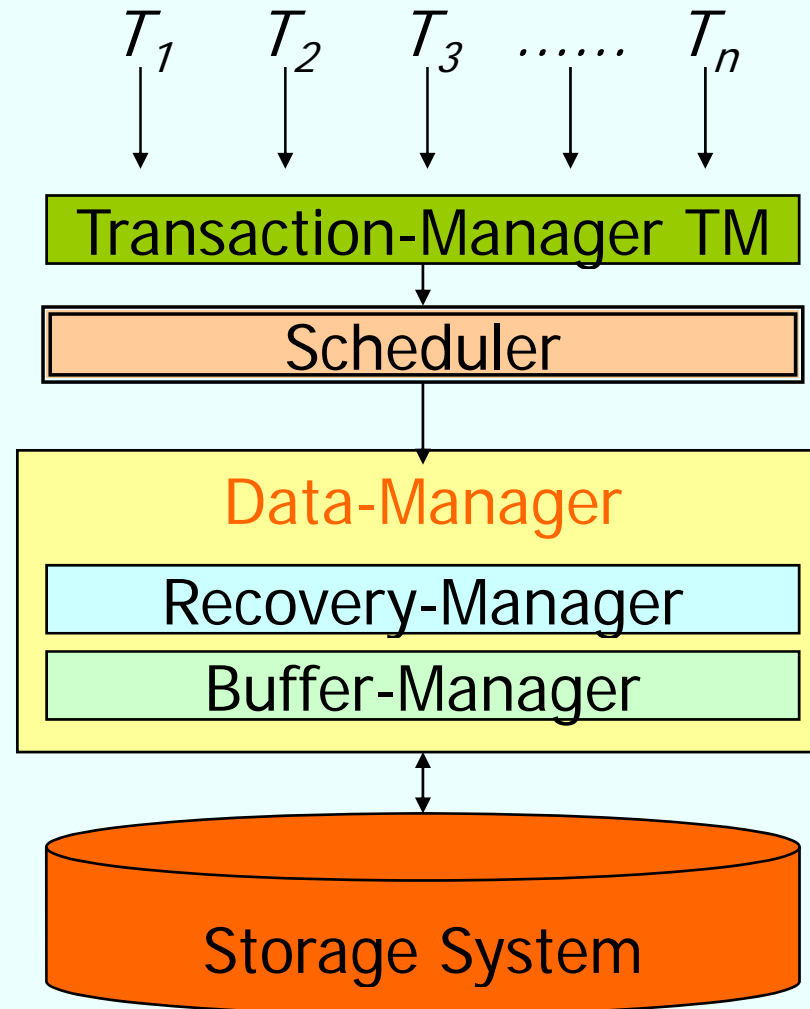
```
values (5275, `Kernphysik`, 3, 2141);
```

insert into Professors

```
values (2141, `Meitner`, `FP`, 205);
```

commit

Database-Scheduler



Concurrency Anomalies

In multi-user operation following concurrency anomalies can occur:

Lost Update

Dirty Read

Non-Repeatable Read

Phantom Reads

Anomalies (2)

Lost Update:

Step	T1	T2
1	read(A, a1)	
2	$a1 = a1 - 300$	
3		read(A, a2)
4		$a2 = a2 * 1,03$
5		write(A, a2)
6	write(A, a1)	
7	read(B, b1)	
8	$b1 = b1 + 300$	
9	write(B, b1)	

t
i
m
e
↓

T1 transfers 300 € from account A to B.

T2 credits account A 3% interest.

Interesting steps:
5 and 6

**update of TA 2
without (again)
reading A overwritten
and thereby lost.**

Anomalies (3)

Dirty Read

Step	T1	T2
1	read(A, a1)	
2	$a1 = a1 - 300$	
3	write(A, a1)	
4		read(A, a2)
5		$a2 = a2 * 1,03$
6		write(A, a2)
7	read(B, b1)	
8	...	
9	abort	

T1 transfers 300 € from account A to B.

T2 credits account A 3% interest.

Interesting steps:
4 and 9

**T1 is aborted,
but T2 has credited
account A the interest in
steps 5/6 - computed
based on the ,wrong'
value of A.**

Anomalies (4)

Non-Repeatable Read

Step	T1	T2
1	select distinct deptnr from emp where salary < 1000	
2		update emp set salary = salary + 10 where deptnr = 2
3	select distinct deptnr from emp where salary < 1000	

T1 lists (twice) all department numbers where there exists an employee with a salary less than 1000.

T2 grants salary increases to all employees from department number 2.

The update of T2 might affect the result of the query in T1.

Anomalies (5)

Phantom Read

Step	T1	T2
1	select sum(balance) from accounts	
2		insert into accounts values (C, 1000)
3	select sum(balance) from accounts	

T1 reads twice the sum of all account balances.

T2 **inserts** a new account with a balance of 1000 €.

T1 computes two different sums.

Synchronization (1)

Criterion for correctness (goal):
logical single usermode, i.e. avoiding all multi
user anomalies

Formal criterion for correctness :
Serializability:

Parallel execution of a set of transactions is
serializable, if there exists **one** serial execution
of the same set of transactions, yielding the

- same data base state and
- the same results as the original execution

Synchronization (2)

But: Serializability restricts parallel execution of transactions

→ accepting anomalies enables less hindrance of transactions
use very **carefully!!**

How to guarantee serializability?

... via *locking*

... via *snapshotting*

Locking (1)

example: RX-locking (simple)

two lock modes:

Read (R)-lock

write- or exclusive (X)-lock

compatibility matrix:

	none	R	X
R	+	+	-
X	+	-	-

"+" means: lock is granted

"-" means: lock conflict

Locking (2)

- with lock conflict requesting transaction has to wait until incompatible lock(s) is (are) removed
- blocking and deadlocks possible
- locks are potentially held until end of transaction

possible optimizations:

- hierarchical locking
- reduced consistency level
- multi version approach

Locking (3)

Incompatibility of a lock request:
→ transaction has to wait

Deadlock:

search for deadlocks in periodical time intervals (adjustable), usually done by cycle detection, resolved by abort of transaction(s)


Timeout: maximum time for waiting for a lock (adjustable), abort of transaction when reached

Deadlock Detection

Wait-for Graph

$T_1 \rightarrow T_2 \rightarrow T_3 \rightarrow T_4 \rightarrow T_1$

$T_2 \rightarrow T_3 \rightarrow T_5 \rightarrow T_2$



- Abort T_3 will resolve both cycles
- Alternative: Deadlock detection with timeouts. Pros/cons?

Consistency levels SQL

four Consistency levels (isolation levels)
determined by the anomalies which may occur
Lost Update always avoided: write locks until end
of transaction

Default: Serializable

	Dirty Read	Non-Repeatable Read	Phantoms
Read Uncommitted	+	+	+
Read Committed	-	+	+
Repeatable Read	-	-	+
Serializable	-	-	-

Consistency levels DB2

	Dirty Read	Non-Repeatable Read	Phantoms
Uncommitted Read (UR)	+	+	+
Cursor Stability (CS)	-	+	+
Read Stability (RS)	-	-	+
Repeatable Read (RR)	-	-	-

Default: Cursor Stability (CS) (!)

Consistency levels PostgreSQL (1)

	Dirty Read	Non-Repeatable Read	Phantoms
Read Uncommitted	+ -	+	+
Read Committed	-	+	+
Repeatable Read	-	-	+ -
Serializable	-	-	-

No anomalies \neq serializable !! (explanation later)

Critique: definition of anomalies stem from a synchronization method using locking

Multiversion concurrency control in PostgreSQL (1)

each transaction sees the database in that state it was when the transaction started
= reads the last committed values that existed at the time it started

called **snapshot isolation**

is a guarantee that all reads made in a transaction will see a consistent snapshot of the database

transaction itself will successfully commit only if no updates it has made conflict with any concurrent updates made since that snapshot

→ only write-write conflicts checked before commit

Multiversion concurrency control in PostgreSQL (2)

- such a write-write conflict will cause the transaction to abort
 - snapshot isolation is implemented by multiversion concurrency control (MVCC)
 - advantage: no reader waits for a writer
no writer waits for a reader
 - disadvantage: needs more space for new versions (no update in place)
needs cleaning
- good if mainly read transactions

Multiversion concurrency control in PostgreSQL (3)

Example: write skew anomaly

T1, T2 start concurrently on the same snapshot

T1 sets V1 to V1 – 200, checks that $V1+V2 \geq 0$

T2 sets V2 to V2 – 200, checks that $V1+V2 \geq 0$

both finally concurrently commit

none has seen the update performed by the other

→ no serializable schedule

but no non-repeatable read anomaly!

snapshot isolation may lead to non serializable schedules

→ serializable snapshot isolation